Experimental Studies on the Value of Information in Financial Markets with Heterogeneously Informed Agents

Jürgen Huber[a], Michael Kirchler[a], Matthias Sutter[b]

[a] Department of Corporate Finance, University of Innsbruck, Austria
[b] Max-Planck-Institute, Jena, Germany

e-mail: michael.kirchler@uibk.ac.at

Abstract

Today information is generally considered the most valuable good in modern economies. Especially in financial markets information is often viewed as the only ingredient necessary to achieve above-average returns. However, empirical, theoretical and experimental work shows that the matter is not that simple. We develop an experimental setting to analyse how valuable forecasting ability is in financial markets. We find that knowledge about the future development of the profits of a company does no necessarily improve the performance of an agent in the market. Our experimental markets show similar behaviour to real markets in several very important aspects, namely volatility clustering, excess kurtosis, and the autocorrelation behaviour. This increases our confidence, that the one feature not observable in real markets – the relation between information level and return – looks similar to our results as well.

JEL-Classification: C91, D82, D83, G14

Keywords: Value of information, financial markets, trading strategies, experimental economics, heterogeneous agents, asymmetric information

1) Introduction

"We live in an information society" is one of the most often heard sentences today. Education, knowledge, and information are considered the most important ingredients to success in business and life. While we agree with this notion for most situations, we think it does not hold for financial markets.

For many years Burton Malkiel (2003a, 2003b) has criticised the underperformance of professional investment funds compared to the index: on average about 70 percent of actively managed stock market funds were outperformed by the market over a ten-year period, for bonds the number is even higher at 90 percent. How can it be that most of the highly paid and professionally trained specialists are not able to do better than the market? We think that our study can give an answer to this question and advice investors how to improve their results: by processing less information.

Beginning with Mandelbrot (1963a, 1963b) we know that returns (price changes) on financial market are not Gaussian distributed, but exhibit excess kurtosis and that these price changes are clustered, so that turbulent and relatively silent phases alternate. Especially in the last years with the emergence of the microsimulation method this field has gained more and more attention. According to Lux/Marchesi (1999, 2000) the clustering effect is mainly based on the appearance of noise in the market, which means that when exceeding some threshold value of noise traders the system becomes unstable and extreme returns occur. But the more agents switch to a noise-strategy the more profitable fundamental trading becomes. So there seems to be a mechanism in the market that lowers volatility until the next phase of turbulences starts. We also test our model on these empirically observed stylized facts to check, whether it is able to capture some of the most well-establised stylized facts observed in real markets.

The innovations of our market are that first, we have a more complex market with agents receiving heterogeneous fundamental information. Second we are using the experimental method to test on stylized facts, which so far has only been done within microsimulations. So we are explaining these empirically observed properties with focus on heterogeneous information in an experimental environment which is a new way of doing this.

2) A market with heterogeneous interacting agents

We develop our model along the lines of Hellwig (1980, 1982), but we are expanding it to 9 information levels. In addition our model does no longer depend on a fundamental value. This is critical, because the fundamental is also not observable in real markets. In most earlier model profit was calculated by comparing prices to the fundamental value – yet, in real markets, only price changes define profits, as the real value is never revealed. Our model follows reality in this respect.

In our market several traders with different forecasting abilities have to trade shares of a virtual company. They can also choose to invest in a risk-free bond. The experimental market is conducted as a continuous double-auction over 30 periods. As one period in the experiment represents one month in real markets, we adjusted several parameters like the risk-free interest rate, the dividends, etc. to monthly data. Each trader was provided with 40 shares (each worth $40 at the start of the experiment) and $1.600 in cash at the start of the experiment. Therefore half of the total wealth was held in cash and the other half in shares. This ensured, that each trader could buy and sell shared depending on his opinion and any endowment effect was prevented. Depending on their expectations about the future development of the share price the participants could then freely buy or sell shares by placing limit orders or accept open bids and asks from other traders.

The most critical part about our market is the information structure: To value the shares agents get information about future dividends. One agent (I1) can only predict this period's dividend correctly. I2 knows the dividends of this and of the next period; I3 knows the actual and the next two dividends, etc. until the 'insider' I9 who knows the actual dividend and who is able to predict the next eight dividends. We therefore have heterogeneously informed agents which are able to predict future cash flows of a company (which is exactly what most analysts try to do, namely calculating the present value of a company by discounting future cash flows). The traders were therefore clearly distinguisehd by the degree of their predicting ability. This allows us to clearly measure the effect of more information on trader behaviour and his return in our market. This is also why we chose the experimental method for our research question: an empirical analysis is not feasible, as it is not possible to distinguish different information levels clearly in real markets – especially the better informed would be hard to find or unwilling to cooperate. A theoretical analysis or a simulation study on the other hand would not be able to capture the complexity of human data processing and the actions it causes in markets.

A major step compared to many other models is, that we have a multi-period model. This means, that traders get new information during the experiment at the start of each new period. At the end of each period the actual dividend is paid out and disappeared from the screen. In turn time t moves one period forward and the dividend of former t+1 is now the actual dividend, while former t+2 is the new t+1, etc. In this way there were 30 periods of trading in the market.

Example of the movement of dividends in two subsequent periods:[1]

	t	t+1	t+2	t+3	t+4	t+5	t+6	t+7	t+8
Dividends shown in period X	20	22	26	23	25	22	18	19	21
Dividends shown in period X+1	22	26	23	25	22	18	19	21	24

The experiment was programmed in z-Tree (Fischbacher, 1999), and in this study we present data from seven markets with nine traders each. The experiments took place in the computer laboratories at the University of Innsbruck in June and July 2004 with Austrian business students. At the start of the experiment each trader was assigned a different information level by chance which was kept constant for the whole experiment. Each trader knew his own endowment and the general distribution of information levels.

3) Experimental Results

The main focus of our experiments is the relationship between the information level of the agents and their relative performance in the market. The second focus is to look at soveral stylized facts observed in real markets and to check whether our market displays them as well.

a. Return per Information Level

The graph below shows the median (solid line) and average (dotted line) of the relative net returns in our seven experimental sessions (dots). We see that information does not display the strictly positive effect on returns that is often assumed. Instead we find that traders with information level I5, who are able to predict the actual and four future dividends correctly, do

[1] E.g.: In period X trader with information level I3 knew the dividends of the actual and the next 2 periods, namely 20, 22, 26. In period (X+1) he saw the dividends 22, 26, 23 for this and the following 2 periods, etc.

worst in the market. The worst informed I1 are able to receive the average market return, which is much better, than some better informed levels.

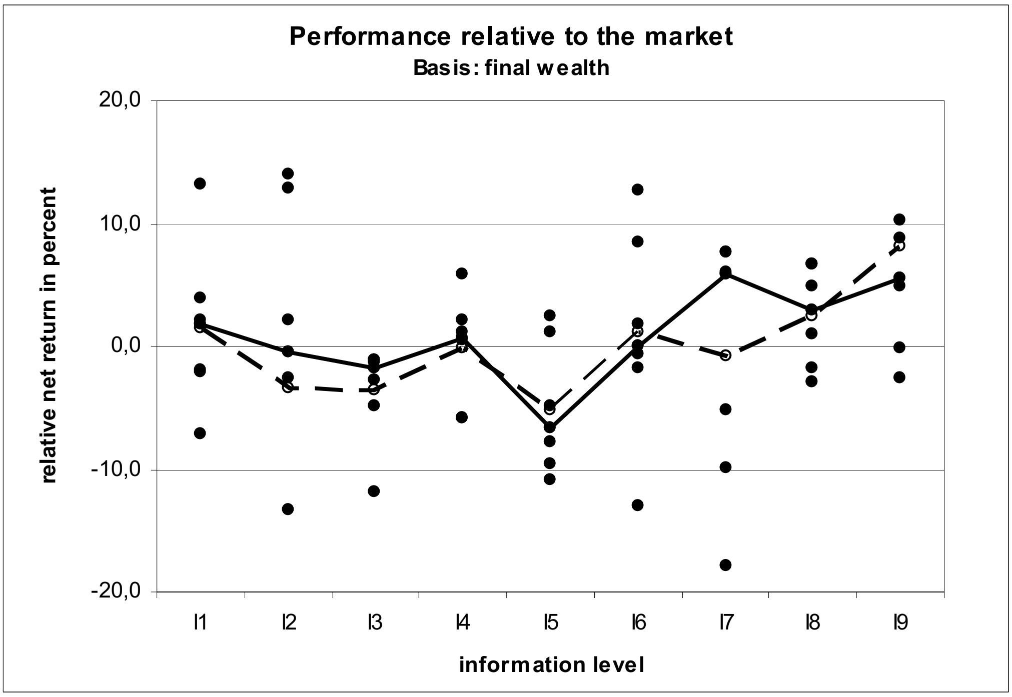

Figure 1: Relative net return in percent per information level

This result contradicts the conventional wisdom, that information is never harmful. While game theory established years ago, that information can be harmful in a strategic game, this was rarely applied to financial markets (Gibbons, 1992). Practitioners are aware, that the financial market is a game where each agent tries to outsmart all others. In this kind of game it is not enough to know a little more that the worst informed – indeed, knowing “something” (e.g. by reading newspapers or watching stock market tv) can be even worse than knowing nothing and just chosing stocks by chance. Insiders will always be able to outperform traders who have just a subset of the information available to them. A trader who does not process information (or just very little information, as I1 does) is behaves unpredictable, as if choosing stocks radomly. In a strategic market context this makes him unexploitable. This is exactly what we see in Figure 1: While the worst informed receive the average market return (which is normalized to zero), insiders are able to gain above-average returns at the cost of the average informed traders.

We want to stress, that this result does not come about to due to inefficient prices. In our analysis of single subject behaviour we found that fundamental information was processed correctly in more than 70 percent of all transactions. Furthermore we see from the movement of prices (red line) and dividends (black line) in Figure 2, that information was mainly processed and reflected in prices. As traders have forecasting abilities prices lead dividends with a lag of five periods. For all markets with exception of market 2 the correlation of prices and dividends [lag=5] is significant at a 1% level (Spearman-Rho: p<0.01).

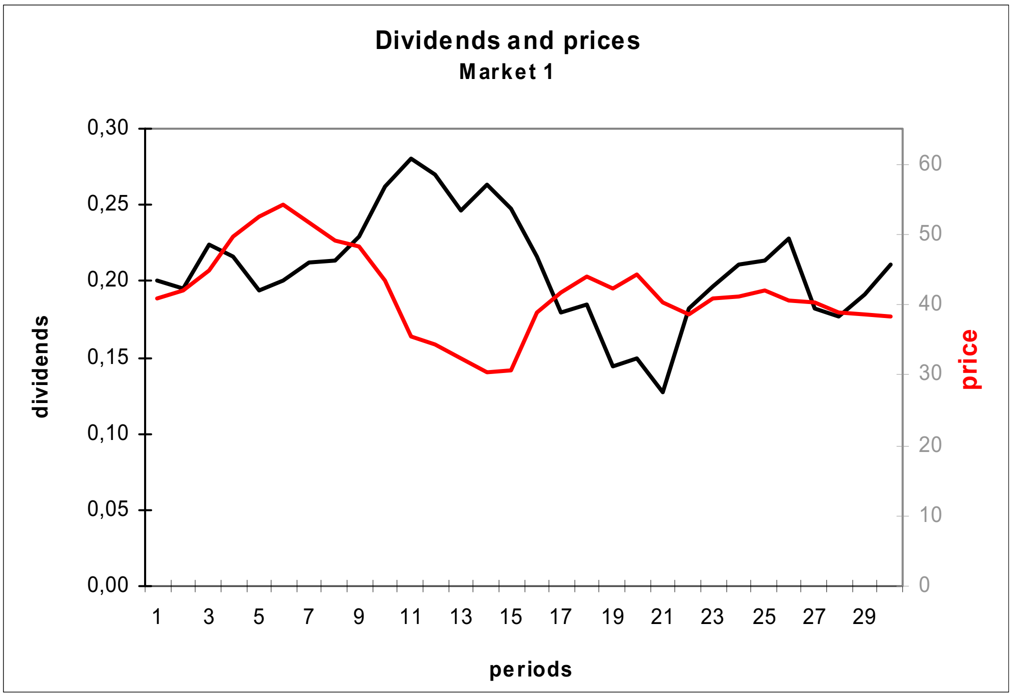

Figure 2: Dividends (black) and mean prices (red) per period in market 1

It is notworthy, that prices lead dividends by about 5 periods, which is exaclty half of the periods that the insider can "predict". The same result – prices refelcting half of the insider information – was also derived by Kyle (1985) with a very different setting in his influential papers. These results are therefore in conflict with Famas (1970) assumption, that prices in a market would reflect all information available.

With our data we can show, that actively acting traders with average information levels will underperform the market, as they buy too late (and therefore too expensive) when prices are rising and vice versa. This is not due to stupidity or errors in information processing, but it is inherent to the market and the information system. These results are very similar to our theoretical predictions, stating that the average informed are the ones that are losing most due to a high covariance in prediction errors with the others. The worst informed traders are somewhat protected by their small amount of information they process, making them basically random traders. So they are more independent than the average informed, because they have a smaller covariance in prediction errors. In a market where the random walk hypothesis holds, this means, that they can expect the average market return. Figure 3 shows the holdings of shares by the different information levels. It can be seen, that first (when prices are still low) the insiders are buying. Two periods later the 'secondary insiders' I7 start buying additional shares. Again two periods later the average informed I5 buy shares, while the insiders (who already observe falling dividends in the future) start selling some of their shares at the now prevailing high prices. With every period worse and worse informed agents start buying shares, but now the prices have already fallen to a lower level.

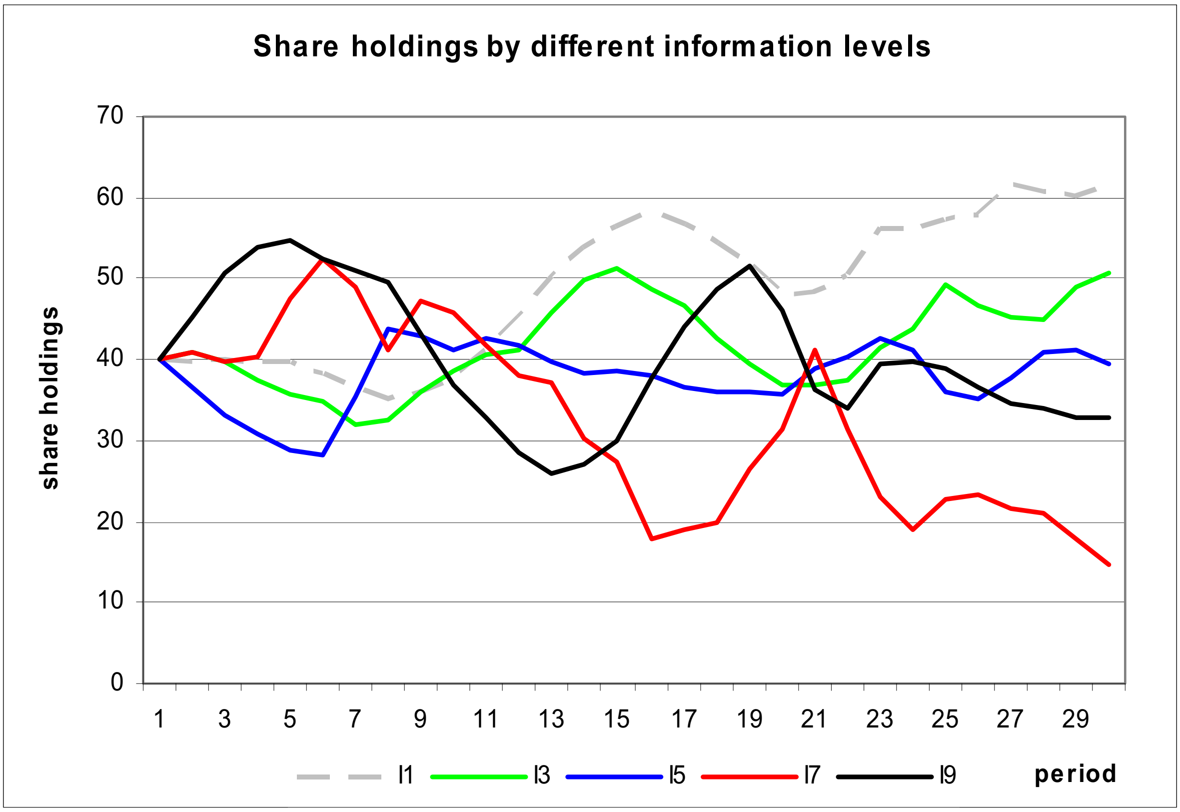

Figure 3: Dividends (black) and mean prices (red) per period in market 1

All of the agents are just acting 'rational', as we really observe, that those information levels who observe relatively the highest (lowest) present values of the asset buy (sell) shares in these periods. Prices also reflect the information given to the traders, although they do not reflect "all available information at any time", but just the average information level. It is astonishing that I5, who has the highest correlation between his estimates and the actual price, and who is actually setting the price very often, is doing worst in the market, when we consider average returns.

b. Empirical Properties of the Market

We tried to construct the main features of our market as realistic as possible to be able to make some reasonable assumptions about the relationship or information level and return in real markets as well. To test not only the 'input' of our assumptions but also the 'output' in form of trading data we conducted some econometric analysis with out experimental data.

Our objective was to test whether our market displays some of the properties of real stock markets, know as stylized facts. First we looked at the distribution of log returns. In most capital markets this distribution displays excess kurtosis, what is known as "fat tails and steep peak".

We found excess kurtosis in our experimental markets with values ranging from 5.47 to 34.46 and non-gausseanity in log returns in all markets according to the Kolmogorov-Smirnov test statistics ($p<0.001$). This can be shown by distributing the log returns. In Figure 3 Markets 5 and 7 are presented. The results for other markets look approximately the same.

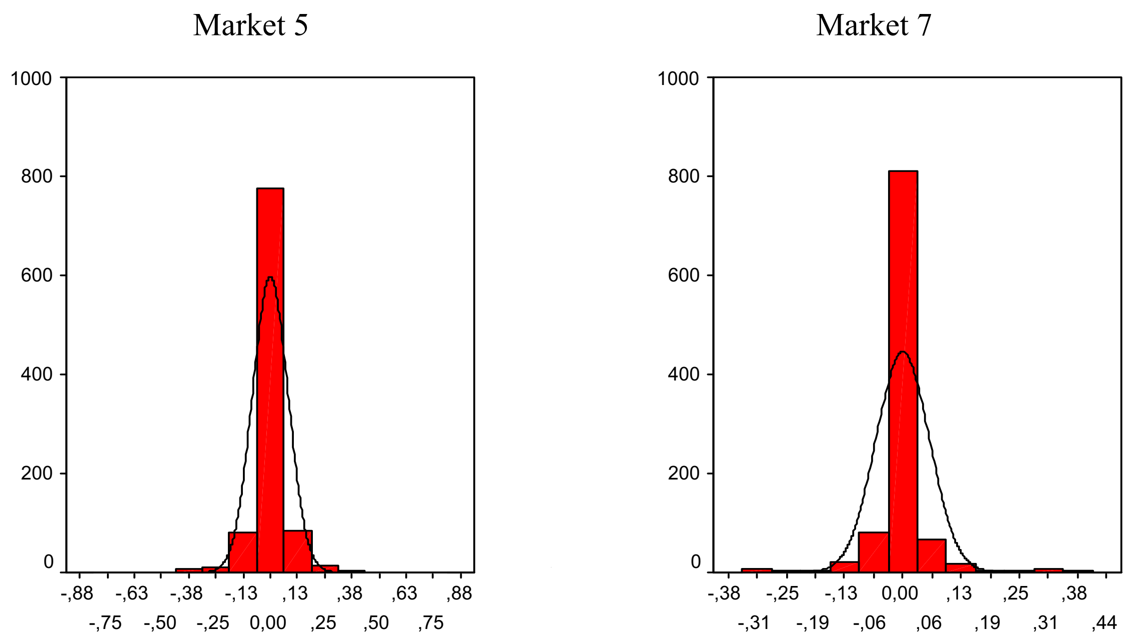

Figure 4: Distribution of log returns of market 5 and market 7

Next we looked at the autocorrelation of log returns. Here we found that the autocorrelation function of price increments is rapidly moving to zero for all markets as is the fact in real financial markets. The high negative autocorelation for lag 1 can mainly be attributed to the trading mechanism, as it is usual for double-auction markets.

Market 5

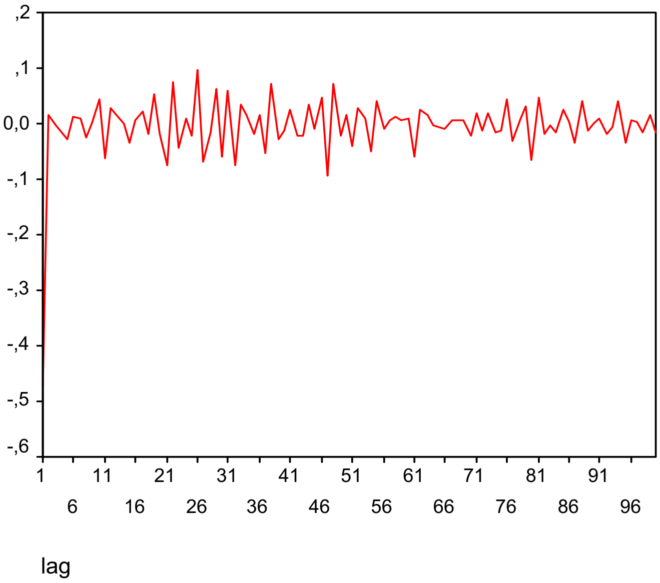

Market 1

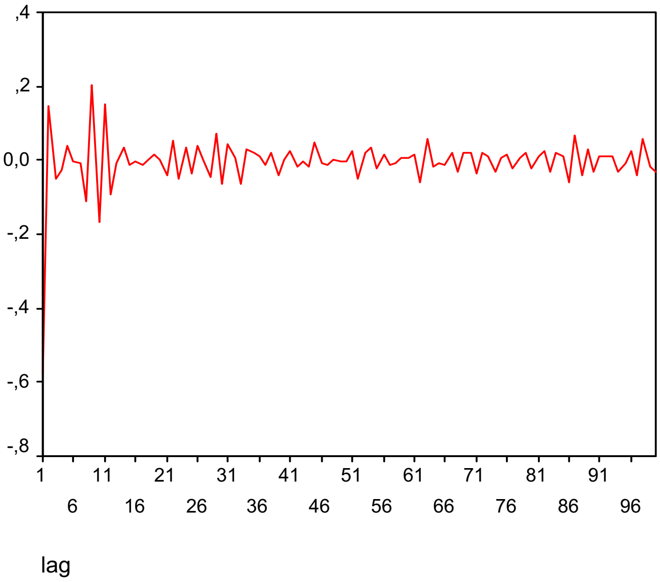

Figure 5: Autocorrelation of log return increments of market 5 and market 1

We also saw that a nonlinear transformation of raw returns, namely absolute returns, exhibit positive autocorrelation or persistence. This is also observed in real markets.

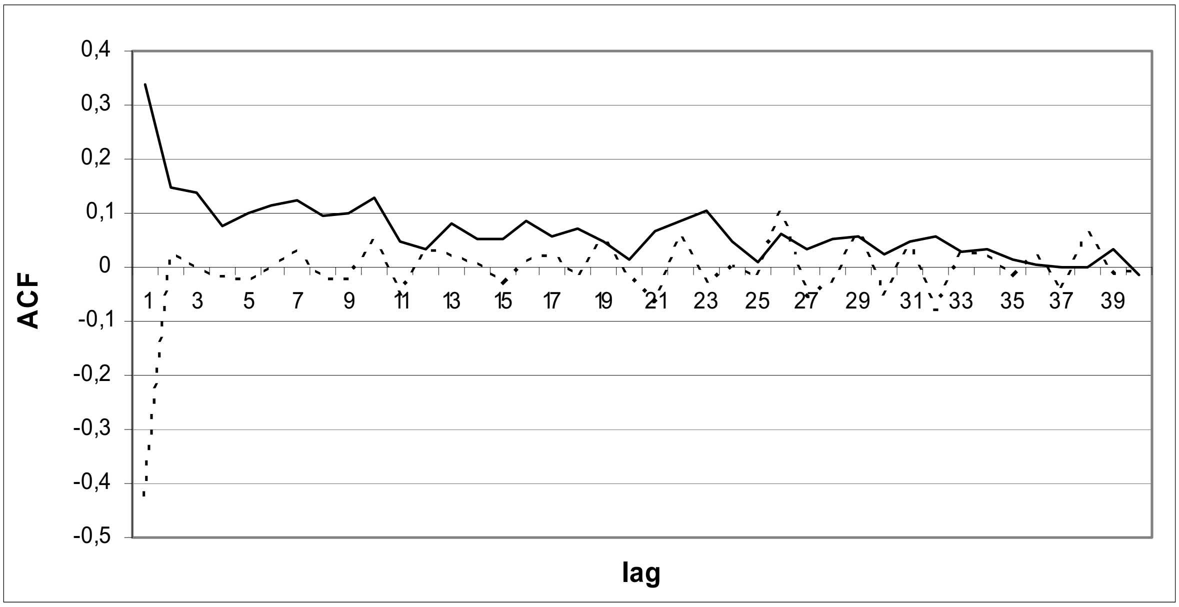

Figure 6: Autocorrelation function of raw returns (dotted line) and absolute returns (solid line) of Market 5. The persistence in autocorrelation of absolute returns is a quantitative signature of the phenomenon of volatility clustering. Large price variations are more likely to be followed by large price variations.

Concerning volatility clustering we can see similar patterns as in real world markets, because turbulent phases and silent phases are alternating.

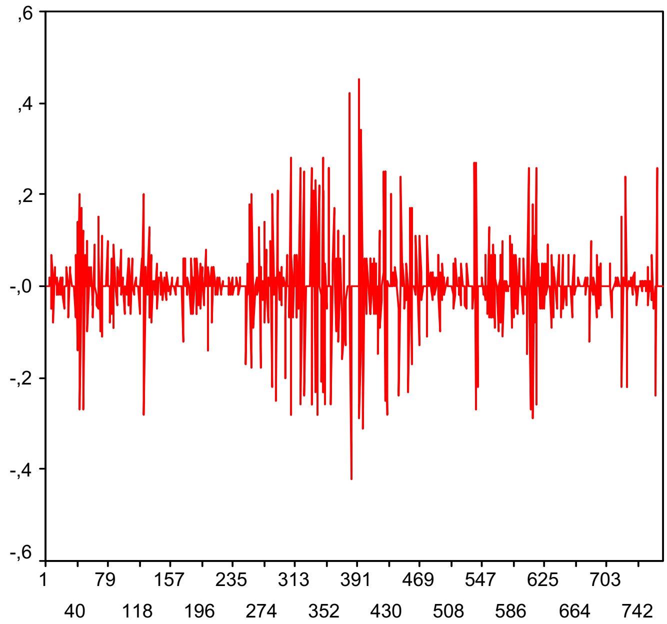

Figure 7: Volatility clustering of log returns of market 6

We found that our markets show similar behaviour to real markets in several very important aspects, namely volatility clustering, excess kurtosis, and the autocorrelation behaviour. This increases our confidence, that the one feature not observable in real markets – the relation between information level and return – looks similar as well to the results we found.

4) Concluding remarks and future work

In this study the relationship between information level and return in a market was explored with experimental data. We understand financial markets as a game, where heterogeneous agents interact. In their attempts to outsmart each other only the best informed are able to gain above average returns. Average informed traders who rely on their information can be exploited by their better informed opponents, while the worst informed are hardly exploitable because they are trading on basis of few information and are therefore close to random traders. So it may be better for agents who are not insiders to stop gathering and processing information.

On the basis of these first motivating results we are looking forward to further testing of the relationship of information level and returns and of the empirical properties of financial markets. In detail we are looking closer to the variation of several variables in our model (experience of traders, trading mechanism, properties of information, etc.) to see their effects on our two major questions, the return distribution and the properties of the markte. There is no work that is dealing with the occurrence and explanation of stylized facts in an experimental environment with heterogeneous fundamentally informed traders so far. Therefore we want to conduct further experiments to find out why the empirically observable stylized facts also emerge in our markets. In our experiments so far we learned that the fat-tail distribution and the effect of volatility clustering have something to do with heterogeneous fundamental information.

References

Bouchaud, J.P., Potters, M., 2001. Theory of Financial Risks. From Statistical Physics to Risk Management. Cambridge, Cambridge University Press.

Cont, R., 1997. Scaling and Correlation in Financial Data. Working Paper

Cont, R., 2001. Empirical Properties of Asset Returns: Stylized Facts and Statistical Issues. Quantitative Finance, Vol. 1, p. 223-236.

Fama, E., 1970. Efficient Capital Markets: A Review of Theory and Empirical Work. Journal of Finance, Vol. 45, p. 383-417.

Fischbacher, U., 1999. Z-tree: Zürich Toolbox for Readymade Economic Experiments. Working paper No. 21, Institute for Empirical Research in Economics, University of Zurich.

Gibbons, R., 1992. A Primer in Game Theory. Prentice Hall. Gosport.

Grossman, S.J., 1976. On the Efficiency of Competitive Stock Markets where Traders have Diverse Information. Journal of Finance, Vol. 51, p. 573-585.

Grossman, S.J., Stiglitz, J.E., 1980. On the Impossibility of Informationally Efficient Prices. American Economic Review, Vol. 70, p. 393-408.

Hellwig, M., 1982. Rational Expectation Equilibrium with Conditioning on Past Prices: A Mean-Variance Example. Journal of Economic Theory, Vol. 26, p. 279-312.

Kyle, A., 1985. Continuous Auctions and Insider Trading. Econometrica 53, p. 1315-1335.

LiCalzi, M., Pellizzari, P., 2003. Fundamentalists Clashing over the Book: A Study of Order-Driven Stock Markets. Quantitative Finance, Vol. 3, p. 1-11.

Lux, T., Marchesi, M., 2000. Volatility Clustering in Financial Markets: A Microsimulation of Interacting Agents. International Journal of Theoretical and Applied Finance, Vol. 3, p. 675-702.

Lux T., Marchesi M., 1999. Scaling and Criticality in a Stochastic Multi-Agent Model of a Financial Market. Letters to Nature, Vol. 397, p. 498-500.

Malkiel, B.G., 2003a. Passive Investment Strategies and Efficient Markets. European Financial Management, Vol. 9 (1), p. 1-10.

Malkiel, B.G., 2003b. The Efficient Market Hypothesis and Its Critics. Journal of Economic Perspectives, Vol. 17 (1), p. 59-82.

Mandelbrot, B., 1963a. New Methods of Statistical Economics. The Journal of Political Economy, Vol. 71, p. 421-440.

Mandelbrot, B., 1963b. The Variation of Certain Speculative Prices. The Journal of Business, Vol. 36, p. 394-419.

Plerou, V., 1999. Scaling of the Distribution of Price Fluctuations of Individual Companies. Physical Review E, p. 6519-6529.

Raberto, M. et.al., 2001. Agent-Based Simulation of a financial market. Physica A, Vol. 299, p. 319-327.

Voit, J., 2003. The Statistical Mechanics of Financial Markets. Berlin, Heidelberg, New York. Springer.